Robot Location Estimation in the Situation Calculus

Vaishak Belle and Hector J. Levesque
Dept. of Computer Science
University of Toronto

ICAPS Workshop on Planning and Robotics, June 2013

Overview

- Motivation
- Formal preliminaries
- Example action theory
- Conclusions
- Future work

Motivation

The situation calculus is a general and rich formalism for representing dynamic worlds:

- serves as foundation for many planning languages
- methodologies such as execution monitoring and loopy plans
- while first-order, practical systems may impose restrictions as they see fit

In the real world, however, effectors and sensors typically noisy

- techniques such as Kalman filtering do indeed address belief propagation in these contexts
- but, very little is said about how actions might change values of certain state variables while not affecting others
- difficult to model strict uncertainty, complex actions that shift dependencies between variables, etc.

Towards a specification

So it becomes imperative that the underlying action formalism, at least in terms of a *specification*, cope with the problems of how the robot is to modify its beliefs based on the actions performed and the results returned by its sensors, even when they are noisy. (IJCAI 2013)

This talk is about demonstrating how such a specification could be used for a delivery robot operating on a planar surface.

- Our setup is a simple one to highlight some of the features that a full formal account of the domain can effectuate.
- We focus on belief change about the robot's location, but we imagine that the robot is manipulating objects, etc.
- Computational considerations discussed as part of future work.

An example to demonstrate features

Robot moving towards the wall: at distance h to it, equipped with a sonar aimed at wall:

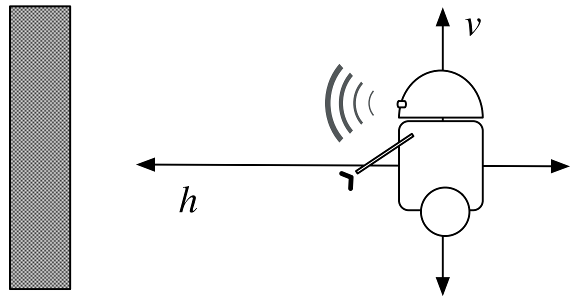

- suppose robot believes h is uniformly distributed on the interval $[2, 12]$
- move by 1 unit (leftwards) shifts distribution on $[1, 11]$
- move by 4 units more radical: $h = 0$ has a *weight* of .2!
- $h \in (0, 8]$ still associated with densities. Mixed distribution retained on a subsequent rightward motion.

An example to demonstrate features (2)

- Assume sonar has additive Gaussian noise. After a sonar reading, beliefs about h's true value should be revised to an appropriate Gaussian.
- Assume a second sensor, say, a GPS device that gives readings for both h and v. Suppose GPS also has a Gaussian error profile, and has systematic bias due to signal obstructions when close to the wall.

  Robot now obtains competing, perhaps conflicting, readings from sonar and GPS about h. How should the robot adjust its beliefs?

Our account handles difficult combinations of continuous sensors, discrete probabilities, probability densities, and shifting dependencies and distributions. It seamlessly integrates logic (strict uncertainty, quantification).

Background: standard situation calculus (Reiter 2001)

- Fluents, situations, actions and objects.
- Situations are *histories*, e.g., $do(a, s)$ unique successor situation of s. Situations can be structured as trees.
- A set of initial situations describes the way the world is initially. S_0 is the actual initial state. Use ι to range over initial states only.

Arrange physical laws in terms of a *basic action theory* $\mathcal{D}$ consisting of

- $\mathcal{D}_0$, which describes what is true initially (any first-order theory);
- preconditions axioms and successor state axioms (incorporating Reiter's monotonic solution to frame problem).

Agents reason by means of entailments of $\mathcal{D}$, e.g.
$\mathcal{D} \models Broken(obj5, do(drop(obj5), do(pickup(obj5), S_0)))$.

Background: continuous uncertainty

We generalize the Bacchus, Halpern and Levesque (BHL) scheme for reasoning about degrees of belief to continuous domains.

Essentials: 2 new distinguished symbols, p and l

- l captures *likelihood* (written like *Poss*):

  $$l(sonar(z), s) = \mathcal{N}(z - h(s); \mu, \sigma^2)$$

  i.e., difference between reading and true value is normally distributed. (In BHL, these are understood to be discrete approximations.)

- p determines a probability distribution on situations: $p(s', s)$ denotes the relative *weight* accorded to situation s' when the agent happens to be in situation s. Initial properties of p specified by modeler as part of $\mathcal{D}_0$. (Example later.)

Background: continuous uncertainty (2)

Framework has only 3 new axioms as part of $\mathcal{D}$:

- only initial situations "epistemically" related to each other, and have nonnegative p values
- a successor state axiom for p, which determines the p value of a successor situation after actions. Roughly, $p(do(a, s'), do(a, s)) = p(s', s) \times l(a, s')$
- *initially*, there is one situation for every vector of fluent values

Then, degree of belief in ϕ is an *abbreviation* for:

$$Bel(\phi, s) \doteq \frac{1}{\gamma} \int_{\vec{x}} \sum_{\vec{y}} Density(\vec{x} \cdot \vec{y}, \phi, s)$$

Here $\vec{x}$ are the initial values of continuous fluents $f_1, \ldots, f_n$, and $\vec{y}$ are the initial values of discrete fluents $g_1, \ldots, g_m$. *Density* is the p value of a situation where ϕ holds, and whose root satisfies $\bigwedge f_i = x_i \wedge \bigwedge g_j = y_j$.

Robot location estimation: Basic action theory

Example action theory $\mathcal{D}$ consists of the three new axioms, which are domain-independent, along with the following sentences.

- h is uniformly distributed, and independently, v is normally distributed:

$$p(\iota, S_0) = \begin{cases} .1 \times \mathcal{N}(v(\iota); 0, 16) & \text{if } 2 \leq h(\iota) \leq 12 \\ 0 & \text{otherwise} \end{cases}$$

- *left* moves robot leftwards (but until the robot hits the wall) and *up* moves it along the Y-axis away from the origin:

$$h(do(a,s)) = u \equiv \exists z(a = left(z) \wedge u = \max(0, h(s) - z)) \vee \neg\exists z(a = left(z)) \wedge u = h(s).$$

$$v(do(a,s)) = u \equiv \exists z(a = up(z) \wedge u = v(s) + z) \vee \neg\exists z(a = up(z)) \wedge u = v(s).$$

Specification of basic action theory $\mathcal{D}$ continued

Two sensors: sonar and GPS, both of which are noisy. Sonar's error prfile

$$l(sonar(z), s) = \mathcal{N}(h(s) - z; 0, .25).$$

Mean 0 indicates no systematic bias. The error profile for the GPS is provided analogously, with systematic bias when the robot is close to the wall. We let the variance in GPS readings be 1, and therefore it is less accurate than the sonar (variance = .25).

This completes the specification of $\mathcal{D}$. We now discuss some entailments.

- $Bel(h = 2 \lor h = 3 \lor h = 4, S_0) = 0$ *initial beliefs*

  Intuitively, although we are integrating a density function $q(x_1, x_2)$ over all real values, $q(x_1, x_2) = 0$ unless $x_1 \in \{2, 3, 4\}$.

- $Bel(5 \leq h \leq 5.5, S_0) = .05$

Logical entailments of $\mathcal{D}$

- $Bel(h = 0, do(left(4), S_0)) = .2$ *physical actions*

  A *continuous* distribution evolves into a *mixed* one. By h's successor state axiom, $h = 0$ holds after the action iff $h \leq 4$ held before.

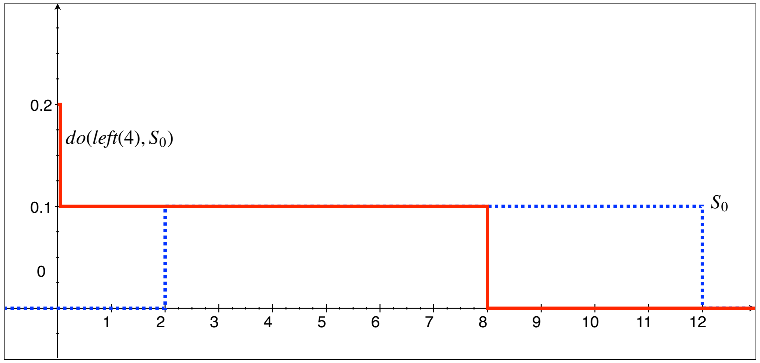

- $Bel(h = 4, do(left(-4), do(left(4), S_0))) = .2$

  $Bel(h = 4, do(left(4), do(left(-4), S_0))) = 0$

  If the robot now moves away, the point $h = 0$ continues to have .2 weight (and obtains a h value of 4). But if the robot had moved away first before moving towards the wall, the distribution remains fully continuous.

Logical entailments of $\mathcal{D}$ (2)

- $Bel(v \leq 1, do(left(6), S_0)) = Bel(v \leq 1, S_0) = \int_{-\infty}^{1} \mathcal{N}(x_2; 0, 16)\mathrm{d}x_2$

  Owing to Reiter's solution to the frame problem, belief in v is unaffected by a lateral motion (which only affects h).

- $Bel(5 \leq h \leq 5.5, do(sonar(5.3), S_0)) \approx .38$ *sonar*

  $Bel(4.5 \leq h \leq 6.5, do(sonar(5.6), do(sonar(5.3), S_0)) \approx .99$

  A single reading sharpens belief, and two successive readings sharpen belief further. Here, readings multiply the p value by sonar's likelihood.

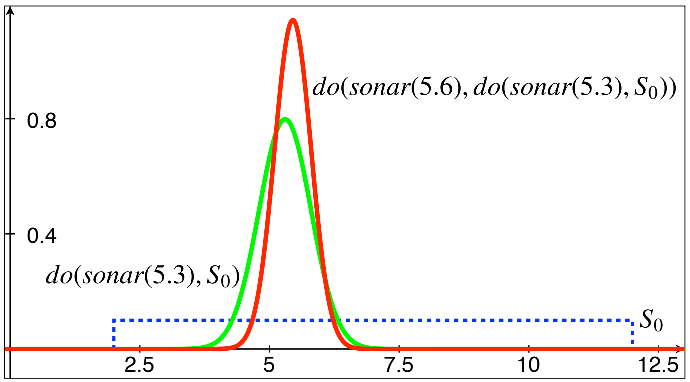

Logical entailments of $\mathcal{D}$ (3)

- $Bel(-1 \le v \le 1, do(gps(5, .1), S_0)) \approx .27$ *GPS*

  The GPS senses both h and v. (Since v has a Gaussian prior, the effect of GPS reading results in another Gaussian, as in Kalman filtering.)

- $Bel(5 \le h \le 5.5, do(gps(5.3, .1), do(gps(5, .1), S_0))) \approx .27$

  $Bel(5 \le h \le 5.5, do(sonar(5.3), do(gps(5, .1), S_0))) \approx .42$

  Sonar is *more sensitive* (lower variance) than the GPS. Its reading is more effective.

Other entailments shown in paper include

- nonstandard properties, e.g., relationships between variables such as $Bel(h > 7v, S_0)$
- reasoning about the past, systematic bias, etc.

Conclusions

- Location estimation for a robot operating in an incompletely known world with noisy sensors.
- Situation calculus + BHL generalization = realistic continuous error models.
- In contrast to a number of competing formalisms, where the modeler is left with the difficult task of deciding how the dependencies and distributions of state variables might evolve, here one need only specify the initial beliefs and the physical laws. Suitable posteriors are then entailed.
- We demonstrated that belief changes appropriately even when one is interested in nonstandard properties and in the presence of actions that affect variables in nontrivial ways, all of which emerges as a side-effect of the general specification.

Future work

Immediate question: a general procedure to effectively reason about beliefs

- One approach: reduce beliefs to what is known initially, i.e. regression

$$\mathcal{D} \models Bel(\phi, do([a_1, \ldots, a_k], S_0)) \text{ iff } \underline{\mathcal{D}_0 \models Bel(\psi, S_0)}$$

  That is, can belief state evolution, including information gained as a result of noisy sensing, be reduced to questions about the initial state?

  Yes! See UAI-13.

- Can we also formally categorize action types that would lead to efficient reasoning?

More broadly, we are interested in the achievability of plans, that is, the question of when can a plan be found and executed, given noisy effector and sensor specifications.